Augmenting Tractable Fragments of Abstract Argumentation*

Sebastian Ordyniak and Stefan Szeider
Vienna University of Technology, Vienna, Austria
ordyniak@kr.tuwien.ac.at, stefan@szeider.net

Abstract

We present a new and compelling approach to the efficient solution of important computational problems that arise in the context of abstract argumentation. Our approach makes known algorithms defined for restricted fragments generally applicable, at a computational cost that scales with the distance from the fragment. Thus, in a certain sense, we gradually augment tractable fragments. Surprisingly, it turns out that some tractable fragments admit such an augmentation and that others do not.

More specifically, we show that the problems of credulous and skeptical acceptance are fixed-parameter tractable when parameterized by the distance from the fragment of acyclic argumentation frameworks. Other tractable fragments such as the fragments of symmetrical and bipartite frameworks seem to prohibit an augmentation: the acceptance problems are already intractable for frameworks at distance 1 from the fragments.

For our study we use a broad setting and consider several different semantics. For the algorithmic results we utilize recent advances in fixed-parameter tractability.

1 Introduction

The study of arguments as abstract entities and their interaction in form of *attacks* as introduced by Dung [1995] has become one of the most active research branches within Artificial Intelligence, Logic and Reasoning [Bench-Capon and Dunne, 2007; Besnard and Hunter, 2008; Rahwan and Simari, 2009]. Abstract argumentation provides suitable concepts and formalisms to study, represent, and process various reasoning problems most prominently in defeasible reasoning (see, e.g., [Pollock, 1992; Bondarenko *et al.*, 1997]) and agent interaction (see, e.g., [Parsons *et al.*, 2003]).

A main issue for any argumentation system is the selection of acceptable sets of arguments, called extensions. However, important computational problems such as determining whether an argument belongs to some extension (Credulous Acceptance) or all extensions (Skeptical Acceptance), are intractable (see, e.g., [Dimopoulos and Torres, 1996; Dunne and Bench-Capon, 2002]). The significance of efficient algorithms for these problems is evident. However, a few tractable fragments are known where the acceptance problems can be efficiently solved: the fragments of acyclic [Dung, 1995], symmetric [Coste-Marquis *et al.*, 2005], bipartite [Dunne, 2007], and noeven [Dunne and Bench-Capon, 2002] argumentation frameworks.

It seems unlikely that an argumentation framework originating from a real-world application belongs to one of the known tractable fragments, but it might be "close" to a tractable fragment. In this paper we study the natural and significant question of whether we can solve the relevant problems efficiently for argumentation frameworks that are of small distance to a tractable fragment. One would certainly have to pay some extra computational cost that increases with the distance from the tractable fragment, but ideally this extra cost should scale gradually with the distance.

*Research funded by the ERC (COMPLEX REASON, 239962).

1.1 Results

We show that the fragments of *acyclic* and *noeven* argumentation frameworks admit an augmentation. In particular, we show that we can solve Credulous and Skeptical Acceptance in polynomial time for argumentation frameworks that are of bounded distance from either of the two fragments. We further show that with respect to the acyclic fragment, the order of the polynomial time bound is independent of the distance, which means that both acceptance problems are *fixed-parameter tractable* (see [Downey and Fellows, 1999]) when parameterized by the distance from the acyclic fragment.

In way of contrast, we show that the fragments of *bipartite* and *symmetric* argumentation frameworks do not admit an augmentation. In particular, we show that the problems Credulous and Skeptical Acceptance are already *intractable* (i.e., (co)NP-hard) for argumentation frameworks at distance 1 from either of the two fragments.

We further show that the parameter "distance to the fragment of acyclic frameworks" is *incomparable* with previously considered parameters that also admit fixed-parameter tractable argumentation [Dunne, 2007; Dvořák *et al.*, 2010]. Hence our approach provides an efficient solution for instances that are hard for known methods.

To get a broad picture of the complexity landscape we take several popular semantics into consideration, namely admissible, preferred, complete, semi-stable and stable semantics (see [Baroni and Giacomin, 2009]).

Our approach is inspired by the notion of "backdoors" which are frequently used in the area of propositional satisfiability (see, e.g., [Williams *et al.*, 2003; Gottlob and Szeider, 2006; Samer and Szeider, 2009]), and also for quantified Boolean formulas and nonmonotonic reasoning [Samer and Szeider, 2009a; Fichte and Szeider, 2011].

2 Preliminaries

An *abstract argumentation system* or *argumentation framework* (*AF*, for short) is a pair (X, A) where X is a finite set of elements called *arguments* and $A \subseteq X \times X$ is a binary relation called *attack relation*. If $(x, y) \in A$ we say that x *attacks* y and that x is an *attacker* of y.

An AF $F = (X, A)$ can be considered as a directed graph, and therefore it is convenient to borrow notions and notation from graph theory. For a set of arguments $Y \subseteq X$ we denote by $F[Y]$ the AF $(Y, \{\, (x, y) \in A \mid x, y \in Y \,\})$ and by $F - Y$ the AF $F[X \setminus Y]$.

Example 1. An AF with arguments $1, \dots, 5$ and attacks $(1, 2)$, $(1, 4)$, $(2, 1)$, $(2, 3)$, $(2, 5)$, $(3, 2)$, $(3, 4)$, $(4, 1)$, $(4, 2)$, $(4, 3)$, $(5, 4)$ is displayed in Fig. 1. □

Let $F = (X, A)$ be an AF, $S \subseteq X$ and $x \in X$. We say that x is *defended* (in F) by S if for each $x' \in X$ such that $(x', x) \in A$ there is an $x'' \in S$ such that $(x'', x') \in A$. We denote by S^+_F the set of arguments $x \in X$ such that either $x \in S$ or there is an $x' \in S$ with $(x', x) \in A$, and we omit the subscript if F is clear from the context. We say S is *conflict-free* if there are no arguments $x, x' \in S$ with $(x, x') \in A$.

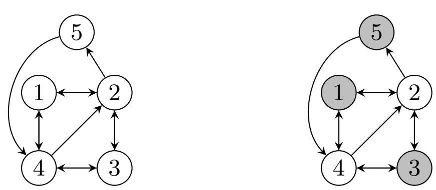

Figure 1: Left: the AF F from Example 1. Right: indicated in gray the only non-empty complete extension of F.

σ	CA_σ	SA_σ
adm	NP-complete	trivial
prf	NP-complete	Π_2^P-complete
com	NP-complete	P-complete
stb	NP-complete	coNP-complete
sem	Σ_2^P-complete	Π_2^P-complete

Table 1: Complexity of credulous and skeptical acceptance for various semantics σ.

Next we define commonly used semantics of AFs, see the survey of Baroni and Giacomin [2009]. We consider a semantics σ as a mapping that assigns to each AF $F = (X, A)$ a family $\sigma(F) \subseteq 2^X$ of sets of arguments, called *extensions*. We denote by adm, prf, com, sem and stb the *admissible*, *preferred*, *complete*, *semi-stable* and *stable* semantics, respectively. These five semantics are characterized by the following conditions which hold for each AF $F = (X, A)$ and each conflict-free set $S \subseteq X$.

- $S \in \mathrm{adm}(F)$ if each $s \in S$ is defended by S.
- $S \in \mathrm{prf}(F)$ if $S \in \mathrm{adm}(F)$ and there is no $T \in \mathrm{adm}(F)$ with $S \subsetneq T$.
- $S \in \mathrm{com}(F)$ if $S \in \mathrm{adm}(F)$ and every argument that is defended by S is contained in S.
- $S \in \mathrm{sem}(F)$ if $S \in \mathrm{adm}(F)$ and there is no $T \in \mathrm{adm}(F)$ with $S^+ \subsetneq T^+$.
- $S \in \mathrm{stb}(F)$ if $S^+ = X$.

Let $F = (X, A)$ be an AF, $x \in X$ and $\sigma \in \{\mathrm{adm}, \mathrm{prf}, \mathrm{com}, \mathrm{sem}, \mathrm{stb}\}$. The argument x is *credulously accepted* in F with respect to σ if x is contained in some extension $S \in \sigma(F)$, and x is *skeptically accepted* in F with respect to σ if x is contained in all extensions $S \in \sigma(F)$.

Each semantics σ gives rise to the following two fundamental computational problems: σ-Credulous Acceptance and σ-Skeptical Acceptance, in symbols CA_σ and SA_σ, respectively. Both problems take as instance an AF $F = (X, A)$ together with an argument $x \in X$. Problem CA_σ asks whether F is credulously accepted in F, problem SA_σ asks whether F is skeptically accepted in F. Table 1, summarizes the complexities of these problems for the considered semantics (see [Dvořák and Woltran, 2010]).

Example 2. Consider the AF F from Example 1 and the complete semantics (com). F has two complete extensions $\emptyset$ and $\{1, 3, 5\}$, see Fig. 1. Consequently, the arguments 1, 3 and 5 are credulously accepted in F and the arguments 2 and 4 are not. Furthermore, because of the complete extension $\emptyset$, no argument of F is skeptically accepted. □

In the following we list classes of AFs for which CA and SA are known to be solvable in polynomial time [Dung, 1995; Baroni and Giacomin, 2009; Coste-Marquis *et al.*, 2005; Dunne, 2007].

- Acyc is the class of *acyclic* argumentation frameworks, i.e., of AFs that do not contain directed cycles.
- Noeven is the class of *noeven* argumentation frameworks, i.e., of AFs that do not contain directed cycles of even length.
- Sym is the class of *symmetric* argumentation frameworks, i.e., of AFs whose attack relation is symmetric.
- Bip is the class of *bipartite* argumentation frameworks, i.e., of AFs whose sets of arguments can be partitioned into two conflict-free sets.

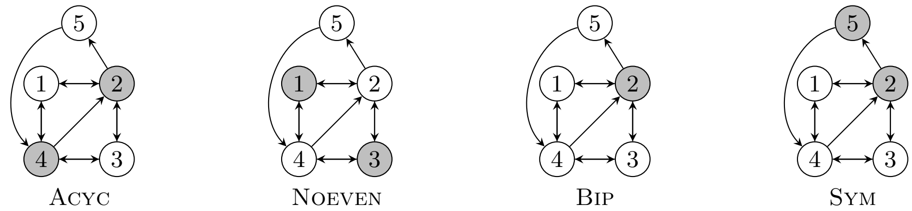

Figure 2: Backdoors for the AF F from Example 1, with respect to the indicated classes.

Lemma 1. *The classes* ACYC, NOEVEN, SYM *and* BIP *can be recognized in polynomial time (i.e., given an AF F, we can decide in polynomial time whether F belongs to any of the four classes).*

Proof. The statement of the lemma is easily seen for the classes ACYC, BIP and SYM. For class NOEVEN it follows by a result of Robertson *et al.* [1999]. □

Since the recognition and the acceptance problems are polynomial for these classes, we consider them as "tractable fragments of abstract argumentation."

2.1 Parameterized Complexity

For our investigation we need to take two measurements into account: the input size n of the given AF F and the distance k of F from a tractable fragment. The theory of *parameterized complexity*, introduced and pioneered by Downey and Fellows [1999], provides the adequate concept and tools for such an investigation. We outline the basic notions of parameterized complexity that are relevant for this paper, for an in-depth treatment we refer to other sources [Flum and Grohe, 2006; Niedermeier, 2006].

An instance of a parameterized problem is a pair (I, k) where I is the *main part* and k is the *parameter*; the latter is usually a non-negative integer. A parameterized problem is *fixed-parameter tractable* (FPT) if there exist a computable function f such that instances (I, k) of size n can be solved in time $f(k) \cdot n^{O(1)}$. Fixed-parameter tractable problems are also called *uniform polynomial-time tractable* because if k is considered constant, then instances with parameter k can be solved in polynomial time where the order of the polynomial is independent of k, in contrast to *non-uniform polynomial-time* running times such as $n^{O(k)}$. Thus we have three complexity categories for parameterized problems: (1) problems that are fixed-parameter tractable (uniform polynomial-time tractable), (2) problems that are non-uniform polynomial-time tractable, and (3) problems that are NP-hard or coNP-hard if the parameter is fixed to some constant (such as k-SAT which is NP-hard for $k = 3$).

2.2 Backdoors

We borrow and adapt the concept of backdoors from the area of propositional satisfiability [Williams *et al.*, 2003; Gottlob and Szeider, 2006; Samer and Szeider, 2009]. Let $\mathcal{C}$ be a class of AFs, $F = (X, A)$ an AF, and $Y \subseteq X$. We call Y a *$\mathcal{C}$-backdoor for F* if $F - Y \in \mathcal{C}$. We write $\text{dist}_{\mathcal{C}}(F)$ for the size of a smallest $\mathcal{C}$-backdoor for F, i.e., $\text{dist}_{\mathcal{C}}(F)$ represents the distance of F from the class $\mathcal{C}$. For an illustration see Fig. 2.

In the following we consider CA and SA parameterized by the distance to a tractable fragment $\mathcal{C}$.

3 Tractability Results

Regarding the fragments of acyclic and noeven argumentation frameworks we obtain the following two results which show that these two fragments admit an amplification.

Theorem 1. *The problems* CA_σ *and* SA_σ *are fixed-parameter tractable for parameter* $\mathrm{dist}_{\textsc{Acyc}}$ *and the semantics* $\sigma \in \{\mathrm{adm}, \mathrm{com}, \mathrm{prf}, \mathrm{sem}, \mathrm{stb}\}$.

Theorem 2. *The problems* CA_σ *and* SA_σ *are solvable in non-uniform polynomial-time for parameter* $\mathrm{dist}_{\textsc{Noeven}}$ *and the semantics* $\sigma \in \{\mathrm{adm}, \mathrm{com}, \mathrm{prf}, \mathrm{sem}, \mathrm{stb}\}$.

The remainder of this section is devoted to a proof of Theorems 1 and 2.

The solution of the acceptance problems involves two tasks: (i) *Backdoor Detection*: to find a $\mathcal{C}$-backdoor B for F of size at most k. (ii) *Backdoor Evaluation*: to use the $\mathcal{C}$-backdoor B for F for deciding whether x is credulously/skeptically accepted in F.

For backdoor detection we utilize recent results from fixed-parameter algorithmics. For backdoor evaluation we introduce and use the new concept of partial labelings.

3.1 Backdoor Detection

The following lemma gives an easy upper bound for the complexity of detecting a $\mathcal{C}$-backdoor for any class $\mathcal{C}$ of AFs that can be recognized in polynomial time.

Proposition 1. *Let* $\mathcal{C}$ *be a class of AFs that can be recognized in polynomial time and* $F = (X, A)$ *an AF with* $\mathrm{dist}_{\mathcal{C}}(F) \leq k$. *Then a* $\mathcal{C}$*-backdoor for* F *of size at most* k *can be found in time* $|X|^{O(k)}$ *and hence in non-uniform polynomial-time for parameter* k.

Proof. To find a $\mathcal{C}$-backdoor for F of size at most k we simply check for every subset $B \subseteq X$ of size $\leq k$ whether $F - B \in \mathcal{C}$. There are $O(|X|^k)$ such sets and each check can be carried out in polynomial time. □

Together with Lemma 1 we obtain the following consequence of Proposition 1.

Corollary 1. *Let* $\mathcal{C} \in \{\textsc{Acyc}, \textsc{Noeven}, \textsc{Sym}, \textsc{Bip}\}$ *and* $F = (X, A)$ *an AF with* $\mathrm{dist}_{\mathcal{C}}(F) \leq k$. *Then a* $\mathcal{C}$*-backdoor for* F *of size at most* k *can be found in time* $|X|^{O(k)}$ *and hence in non-uniform polynomial-time for parameter* k.

It is a natural question to ask whether the above result can be improved to uniform-polynomial time. We get an affirmative answer for three of the four classes under consideration.

Lemma 2. *Let* $\mathcal{C} \in \{\textsc{Acyc}, \textsc{Sym}, \textsc{Bip}\}$ *and* $F = (X, A)$ *an AF with* $\mathrm{dist}_{\mathcal{C}}(F) \leq k$. *Then the detection of a* $\mathcal{C}$*-backdoor for* F *of size at most* k *is fixed-parameter tractable for parameter* k.

Proof. The detection of Acyc-backdoors is easily seen to be equivalent to the so-called directed feedback vertex set problem which has recently been shown to be fixed-parameter tractable by Chen *et al.* [2008]. Similarly, the detection of Bip-backdoors is equivalent to the problem of finding an odd cycle traversal which is fixed-parameter tractable due to a result of Reed *et al.* [2004]. Finally, the detection of a Sym-backdoor set is equivalent to the vertex cover problem which is well known to be fixed-parameter tractable [Downey and Fellows, 1999]. □

We must leave it open whether the detection of Noeven-backdoors of size at most k is fixed-parameter tractable for parameter k. Since already the polynomial-time recognition of Noeven is highly nontrivial, a solution for the backdoor problem seems very challenging. However, it is easy to see that $\mathcal{C}$-backdoor detection, considered as a non-parameterized problem, where k is just a part of the input, is NP-complete for $\mathcal{C} \in \{\textsc{Acyc}, \textsc{Noeven}, \textsc{Sym}, \textsc{Bip}\}$. Hence it is unlikely that Lemma 2 can be improved to a polynomial-time result.

3.2 Backdoor Evaluation

Let $F = (X, A)$ be an AF. A *partial labeling* of F, or *labeling* for short, is a function $\lambda : Y \to \{\text{IN}, \text{OUT}, \text{UND}\}$ defined on a subset Y of X. Partial labelings generalize *total* labelings which are defined on the entire set X of arguments [Modgil and Caminada, 2009].

We denote by $\text{IN}(\lambda)$, $\text{OUT}(\lambda)$ and $\text{UND}(\lambda)$ the sets of arguments $x \in X$ with $\lambda(x) = \text{IN}$, $\lambda(x) = \text{OUT}$ and $\lambda(x) = \text{UND}$ respectively. Furthermore, we set $\text{DEF}(\lambda) = Y$ and $\text{UD}(\lambda) = X \setminus \text{DEF}(\lambda)$ and denote by $\lambda_\emptyset$ the the *empty labeling*, i.e., the labeling with $\text{DEF}(\lambda_\emptyset) = \emptyset$. For a set $S \subseteq X$ we define $\text{lab}(F, S)$ to be the *labeling of F with respect to S* by setting $\text{IN}(\text{lab}(F, S)) = S$, $\text{OUT}(\text{lab}(F, S)) = S^+ \setminus S$ and $\text{UND}(\text{lab}(F, S)) = X \setminus S^+$. We say a set $S \subseteq X$ is *compatible* with a labeling λ if $\lambda(x) = \text{lab}(F, S)(x)$ for every $x \in \text{DEF}(\lambda)$.

Let $F = (X, A)$ be an AF and λ a partial labeling of F. The *propagation* of λ with respect to F, denoted λ^*, is the labeling that is obtained from λ by initially setting $\lambda^*(x) = \lambda(x)$, for every $x \in \text{DEF}(\lambda)$, and subsequently applying one of the following three rules to unlabeled arguments $x \in X$ as long as possible.

Rule 1. x is labeled OUT if x has at least one attacker that is labeled IN.

Rule 2. x is labeled IN if all attackers of x are labeled OUT.

Rule 3. x is labeled UND if all attackers of x are either labeled OUT or UND and at least one attacker of x is labeled UND.

It is easy to see that λ^* is well-defined and unique.

For an AF F, a set B of arguments of F and a partial labeling λ of F we set:

$$\begin{aligned} \text{com}^*(F, \lambda) &= \{\, \text{IN}(\lambda^*) \cup S \mid S \in \text{adm}(F - \text{DEF}(\lambda^*)) \,\}; \\ \text{com}^*(F, B) &= \textstyle\bigcup_{\lambda : B \to \{\text{IN},\text{OUT},\text{UND}\}} \text{com}^*(F, \lambda). \end{aligned}$$

The following lemmas illustrate the connection between partial labelings and complete extensions.

Lemma 3. *Let $F = (X, A)$ be an AF, λ a partial labeling of F, and S a complete extension that is compatible with λ. Then the propagation λ^* of λ is compatible with S.*

Proof. We show the claim by induction on the number of arguments that have been labeled according to Rules 1–3. Because S is compatible with λ it holds that $\lambda^*(x) = \lambda(x) = \text{lab}(F, S)(x)$ for every $x \in \text{DEF}(\lambda)$ and hence the proposition holds before the first argument has been labeled according to one of the rules. Now, suppose that λ' is the labeling that is obtained from λ after labeling the first i arguments according to one of the rules and that x is the $i+1$-th argument that is labeled according to the rules. We distinguish three cases.

First we assume that x is labeled according to Rule 1. In this case $\lambda^*(x) = \text{OUT}$ and we need to show that $x \in S^+ \setminus S$. It follows from the definition of Rule 1 that x has at least one attacker n with $\lambda'(n) = \text{IN}$. Using the induction hypothesis it follows that $\text{lab}(F, S)(n) = \text{IN}$ and hence $n \in S$. Because S is conflict-free it follows that $x \notin S$ but since x is attacked by n it follows that $x \in S^+ \setminus S$.

Second we assume that x is labeled according to Rule 2. In this case $\lambda^*(x) = \text{IN}$ and we need to show that $x \in S$. Let $n_1, \dots, n_r$ be all the attackers of x in F. It follows from the definition of Rule 2 that $\lambda'(n_j) = \text{OUT}$ for every $1 \leq j \leq r$. Using the induction hypothesis it follows that $\text{lab}(F, S)(n_j) = \text{OUT}$ and hence $n_j \in S^+ \setminus S$ for every $1 \leq j \leq r$. It follows that no out-neighbor of x can be contained in S otherwise this out-neighbor would be attacked by x but x cannot be defended by S. Hence $S \cup \{x\}$ is also admissible. Because x is defended by S it follows that $x \in S$.

Finally, we assume that x is labeled according to Rule 3. In this case $\lambda^*(x) = \text{UND}$ and we need to show that $x \notin S^+$. Using the definition of Rule 3 it follows that the set of all attackers of x can be partitioned into two sets U and O such that $\lambda'(u) = \text{UND}$ for every $u \in U$ and $\lambda'(o) = \text{OUT}$ for every $o \in O$ and $U \neq \emptyset$. Using the induction hypothesis it follows that $\lambda'(n) = \text{lab}(F, S)(n)$ for every $n \in U \cup O$. Hence, no attacker of x belongs to S and so x cannot be contained in $S^+ \setminus S$. Furthermore, because S is admissible

and x has an attacker that is not contained in S^+ it follows that x cannot be contained in S. Hence, x is not contained in S^+. $\square$

Lemma 4. *Let $F = (X, A)$ be an AF and $B \subseteq X$. Then* $\mathrm{com}(F) \subseteq \mathrm{com}^*(F, B)$.

Proof. Let $F = (X, A)$ be the given AF, $B \subseteq X$ and $S \in \mathrm{com}(F)$. We show that $S \in \mathrm{com}^*(F, \lambda) = \{\, \textsc{in}(\lambda^*) \cup S \mid S \in \mathrm{adm}(F - \textsc{def}(\lambda^*)) \,\}$ for the unique partial labeling λ defined on B that is compatible with S. We set $S_1 = S \cap \textsc{def}(\lambda^*)$, $S_2 = S \setminus S_1$, and $F_2 = F - \textsc{def}(\lambda^*)$.

It follows from Lemma 3 that $S_1 = \textsc{in}(\lambda^*)$. It remains to show that S_2 is admissible in F_2. Clearly, S_2 is conflict-free. To see that S_2 is admissible suppose to the contrary that there is an argument $x \in S_2$ that is not defended by S_2 in F_2, i.e., x has an attacker y in F_2 that is not attacked by an argument in S_2. Because S is a complete extension of F the argument x is defended by S in F. Hence, there is a $z \in S_1 = S_1 = \textsc{in}(\lambda^*)$ that attacks y. But then, using rule Rule 1, $\lambda^*(y) = \textsc{out}$, and hence y cannot be an argument of F_2. Hence S_2 is admissible in F_2. $\square$

For an AF F we set $F^* = F - \textsc{def}(\lambda^*_{\emptyset})$. In other words, F^* is obtained from F after deleting all arguments from F that, starting from the empty labeling, are labeled according to the Rules 1–3. We observe that because we start from the empty labeling Rule 3 will not be invoked.

We say a class $\mathcal{C}$ of AFs is *fully tractable* if (i) for every $F \in \mathcal{C}$ the set $\mathrm{adm}(F^*)$ can be computed in polynomial time, and (ii) $\mathcal{C}$ is closed under the deletion of arguments, i.e., if $F = (X, A) \in \mathcal{C}$ and $Y \subseteq X$, then $F - Y \in \mathcal{C}$.

Theorem 3. *Let $\mathcal{C}$ be a fully tractable class of AFs, $F = (X, A)$ an AF and B a $\mathcal{C}$-backdoor for F with $|B| \leq k$. Then the computation of the sets* $\mathrm{com}(F)$, $\mathrm{prf}(F)$, $\mathrm{sem}(F)$ *and* $\mathrm{stb}(F)$ *can be carried out in time $3^k|X|^{O(1)}$ and is therefore fixed-parameter tractable for parameter k.*

Proof. Let $\mathcal{C}$ be a fully tractable class of AFs, $F = (X, A)$ an AF and B a $\mathcal{C}$-backdoor for F with $|B| \leq k$. We first show that the computation of $\mathrm{com}(F)$ is fixed-parameter tractable for parameter k. Let λ be one of the 3^k partial labelings of F defined on B. We first show that we can compute $\mathrm{com}^*(F, \lambda) = \{\, \textsc{in}(\lambda^*) \cup S \mid S \in \mathrm{adm}(F - \textsc{def}(\lambda^*)) \,\}$ in polynomial time, i.e., in time $|X|^{O(1)}$. Clearly, we can compute the propagation λ^* of λ in polynomial time. Furthermore, because $F - B \in \mathcal{C}$ (B is a $\mathcal{C}$-backdoor) also $F - \textsc{def}(\lambda^*) \in \mathcal{C}$; this follows since $B \subseteq \textsc{def}(\lambda^*)$ and $\mathcal{C}$ is closed under argument deletion since $\mathcal{C}$ is assumed to be fully tractable. Moreover, since $\mathcal{C}$ is assumed to be fully tractable and $F - \textsc{def}(\lambda^*) \in \mathcal{C}$, we can compute $\mathrm{adm}((F - \textsc{def}(\lambda^*))^*) = \mathrm{adm}(F - \textsc{def}(\lambda^*))$ in polynomial time. Consequently, we can compute the set $\mathrm{com}^*(F, \lambda)$ in polynomial time. Since there are at most 3^k partial labelings of F defined on B, it follows that we can compute the entire set $\mathrm{com}^*(F, B)$ in time $3^k|X|^{O(1)}$.

By Lemma 4 we have $\mathrm{com}(F) \subseteq \mathrm{com}^*(F, B)$. Thus we can obtain $\mathrm{com}(F)$ from $\mathrm{com}^*(F, B)$ by simply testing for each $S \in \mathrm{com}^*(F, B)$ whether S is a complete extension of F. It is a well-known fact that each such a test can be carried out in polynomial time (see e.g. [Dvořák and Woltran, 2010]). Hence, we conclude that indeed $\mathrm{com}(F)$ can be computed in time $3^k|X|^{O(1)}$.

For the remaining sets $\mathrm{prf}(F)$, $\mathrm{sem}(F)$ and $\mathrm{stb}(F)$ we note that each of them is a subset of $\mathrm{com}(F)$. Furthermore, the extensions in $\mathrm{prf}(F)$ are exactly the extensions in $\mathrm{com}(F)$ which are maximal with respect to set inclusion. Similarly, the extensions in $\mathrm{sem}(F)$ are exactly the extensions in $S \in \mathrm{com}(F)$ where the set S^+ is maximal with respect to set inclusion, and $\mathrm{stb}(F)$ are exactly the extensions $S \in \mathrm{com}(F)$ where $S^+ = X$. Clearly, these observations can be turned into an algorithm that computes from $\mathrm{com}(F)$ the sets $\mathrm{prf}(F), \mathrm{sem}(F), \mathrm{stb}(F)$ in polynomial time. $\square$

Lemma 5. *The classes* $\textsc{Acyc}$ *and* $\textsc{Noeven}$ *are fully tractable.*

Proof. It is easy to see that both classes satisfy condition (ii) of being fully tractable, i.e., both classes are closed under the deletion of arguments. It remains to show that they also satisfy condition (i) of being fully tractable, i.e., for every $F \in \textsc{Acyc} \cup \textsc{Noeven} = \textsc{Noeven}$ it holds that the set $\mathrm{adm}(F^*)$ can be

computed in polynomial time. Dunne and Bench-Capon [2001] have shown that if $F \in$ NOEVEN and every argument of F is contained in at least one directed cycle, then $\text{adm}(F) = \{\emptyset\}$. Consequently, it remains to show that if $F \in$ NOEVEN then every argument of F^* lies on a directed cycle. To see this it suffices to show that every argument x of F^* has at least one attacker in F^*. Suppose not, i.e., there is an argument $x \in X \setminus \text{DEF}(\lambda^*_\emptyset)$ with no attacker in F^*. It follows that every attacker of x must be labeled and hence $x \in \text{DEF}(\lambda^*_\emptyset)$, a contradiction. □

Combining Theorem 3 with Lemma 5 we conclude that if $\mathcal{C} \in \{\text{ACYC}, \text{NOEVEN}\}$ then the backdoor evaluation problem is fixed-parameter tractable parameterized by the size of the backdoor set for the semantics $\sigma \in \{\text{com}, \text{prf}, \text{sem}, \text{stb}\}$. For the remaining case of admissible semantics, we recall from Table 1 that SA_{adm} is trivial. Furthermore, we observe that every admissible extension is contained in some complete extension, and every complete extension is also admissible. We conclude that an argument is credulously accepted with respect to the admissible semantics if and only if the argument is credulously accepted with respect to complete semantics. Hence, we have shown that backdoor evaluation is also fixed-parameter tractable with respect to admissible semantics. Together with Lemma 2 and Lemma 1 this establishes our main results Theorem 1 and Theorem 2 of this section.

λ		λ^*				$\text{IN}(\lambda^*) \in$
2	4	1	3	5	$\text{IN}(\lambda^*)$	$\text{com}(F)$?
IN	IN	OUT	OUT	OUT	$\{2,4\}$	no
IN	OUT	OUT	OUT	OUT	$\{2\}$	no
IN	UND	OUT	OUT	OUT	$\{2\}$	no
OUT	IN	OUT	OUT	IN	$\{4,5\}$	no
OUT	OUT	IN	IN	IN	$\{1,3,5\}$	yes
OUT	UND	UND	UND	IN	$\{5\}$	no
UND	IN	OUT	OUT	UND	$\{4\}$	no
UND	OUT	UND	UND	UND	$\emptyset$	yes
UND	UND	UND	UND	UND	$\emptyset$	yes

Table 2: Calculation of all complete extensions for the AF F of Example 1 using the ACYC-backdoor $\{2,4\}$.

Example 3. Consider again the AF F from Example 1. We have observed above that F has an ACYC-backdoor B consisting of the arguments 2 and 4. We now show how to use the backdoor B to compute all complete extensions of F using the procedure given in Theorem 3. Table 2 shows the propagations for all partial labelings of F defined on B together with the set $\text{IN}(\lambda^*)$ and for every λ it is indicated whether the set $\text{IN}(\lambda^*)$ is a complete extension of F. Because $F - B$ is acyclic it follows that $\text{adm}(F - \text{DEF}(\lambda^*)) = \{\emptyset\}$ (see the proof of Lemma 5) and hence $\text{com}^*(F, \lambda) = \{\text{IN}(\lambda^*)\}$. It is now easy to compute $\text{com}^*(F, B)$ as the union of all the sets $\text{IN}(\lambda^*)$ given in Table 2. Furthermore, using the rightmost column of Table 2 we conclude that $\text{com}(F) = \{\emptyset, \{1,3,5\}\}$, which is in full correspondence to our original observation in Example 2. □

4 Hardness Results

The hardness results for CA_σ and SA_σ are not completely symmetric since for $\sigma \in \{\text{adm}, \text{com}\}$ the former problem is NP-complete, the latter is solvable in polynomial time (recall Table 1).

Theorem 4. *(1) The problem* CA_σ *is* NP*-hard for AFs F with* $\text{dist}_{\text{BIP}}(F) = 1$ *and* $\sigma \in \{\text{adm, com, prf, sem, stb}\}$. *(2) The problem* SA_σ *is* coNP*-hard for AFs F with* $\text{dist}_{\text{BIP}}(F) = 1$ *and* $\sigma \in \{\text{prf}, \text{sem}, \text{stb}\}$.

Proof. (Sketch.) The hardness results follow by reductions from Monotone 3-Satisfiability [Garey and Johnson, 1979] and its complement, similar to reductions used by Dunne [2007]. We illustrate the constructions in Fig. 3. $\square$

Figure 3: Illustrations for the reductions in the proof of Theorem 4, showing instances (F, φ) and (F', φ') for the problems CA_σ and SA_σ, respectively, obtained from the monotone 3-CNF formula $\varphi = C_1 \wedge \overline{C}_1$ with $C_1 = x_1 \vee x_2 \vee x_3$ and $\overline{C}_1 = \neg x_1 \vee \neg x_2 \vee \neg x_3$. The set $\{\varphi\}$ is a Bip-backdoor for F and F'.

Theorem 5. *(1) The problem* CA_σ *is* NP*-hard for AFs F with* $\mathrm{dist}_{\mathrm{SYM}}(F) = 1$ *and* $\sigma \in \{\mathrm{adm}, \mathrm{com}, \mathrm{prf}, \mathrm{sem}, \mathrm{stb}\}$. *(2) The problem* SA_σ *is* coNP*-hard for AFs F with* $\mathrm{dist}_{\mathrm{SYM}}(F) = 1$ *and* $\sigma \in \{\mathrm{prf}, \mathrm{sem}, \mathrm{stb}\}$.

Proof. (Sketch.) We use a reduction from 3-Satisfiability [Garey and Johnson, 1979] and its complementary problem, similar to reductions used by Dimopoulos and Torres [1996]. We illustrate the reductions in Fig. 4. $\square$

Figure 4: Illustrations for the reductions in proof of Theorem 5. (F, φ) and (F', φ') are instances of CA_σ and SA_σ, respectively, obtained from the 3-CNF formula $\varphi = C_1 \wedge C_2 \wedge C_3$ with $C_1 = x_1 \vee x_2 \vee x_3$, $C_2 = \neg x_1 \vee x_2 \vee \neg x_3$ and $C_3 = \neg x_1 \vee \neg x_2 \vee \neg x_3$. The set $\{\varphi\}$ is a Sym-backdoor for F and F'.

5 Comparison with other Parameters

In this section we compare our new structural parameters $\mathrm{dist}_{\mathrm{ACYC}}$ and $\mathrm{dist}_{\mathrm{NOEVEN}}$ to the parameters treewidth and clique-width that have been introduced to the field of abstract argumentation by Dunne [2007] and Dvořák *et al.* [2010], respectively. Due to space requirements we cannot give the definitions of these parameters and must refer the reader to the above references. The following two propositions show that treewidth and clique-width are both incomparable to our distance parameters.

Proposition 2. *There are acyclic and noeven AFs that have arbitrarily high treewidth and clique-width.*

Proof. Consider any symmetric AF F of high treewidth or clique-width together with an arbitrary but fixed ordering $<$ of the arguments of F. By deleting all attacks from an argument x to an argument y with $y < x$ we obtain an acyclic AF F' whose treewidth and clique-width is at least as high as the treewidth and clique-width of F, but $\text{dist}_{\text{NOEVEN}}(F) = \text{dist}_{\text{ACYC}}(F) = 0$. □

Proposition 3. *There are AFs with bounded treewidth and clique-width where* $\text{dist}_{\text{NOEVEN}}$ *and* $\text{dist}_{\text{ACYC}}$ *are arbitrarily high.*

Proof. Consider the AF F that consists of n disjoint directed cycles of even length. It is easy to see that the treewidth and the clique-width of F are bounded by a constant but $\text{dist}_{\text{NOEVEN}}(F) = \text{dist}_{\text{ACYC}}(F) = n$. □

6 Conclusion

We have introduced a novel approach to the efficient solution of acceptance problems for abstract argumentation frameworks by "augmenting" a tractable fragment. This way the efficient solving techniques known for a restricted fragment, like the fragment of acyclic argumentation frameworks, become generally applicable to a wider range of argumentation frameworks and thus relevant for real-world instances. Our approach is orthogonal to decomposition-based approaches and thus we can solve instances efficiently that are hard for known methods.

The augmentation approach entails two tasks, the detection of a small backdoor and the evaluation of the backdoor. For the first task we could utilize recent results from fixed-parameter algorithm design, thus making results from a different research field applicable to abstract argumentation. For the second task we have introduced the concept of partial labelings, which seems to us a useful tool that may be of independent interest. In view of the possibility of an augmentation, our results add significance to known tractable fragments and motivate the identification of new tractable fragments. For future research we plan to extend our results to other semantics and new tractable fragments.

References

[Baroni and Giacomin, 2009] P. Baroni and M. Giacomin. Semantics of abstract argument systems. In I. Rahwan and G. Simari, eds, *Argumentation in Artificial Intelligence*, pages 25–44. Springer Verlag, 2009.

[Bench-Capon and Dunne, 2007] T. J. M. Bench-Capon and P. E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

[Besnard and Hunter, 2008] P. Besnard and A. Hunter. *Elements of Argumentation.* The MIT Press, 2008.

[Bondarenko *et al.*, 1997] A. Bondarenko, P. M. Dung, R. A. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93(1-2):63–101, 1997.

[Chen *et al.*, 2008] J. Chen, Y. Liu, S. Lu, B. O'Sullivan, and I. Razgon. A fixed-parameter algorithm for the directed feedback vertex set problem. *J. ACM*, 55(5):Art. 21, 19, 2008.

[Coste-Marquis *et al.*, 2005] S. Coste-Marquis, C. Devred, and P. Marquis. Symmetric argumentation frameworks. In L. Godo, ed., *ECSQARU 2005*, *LNCS* 3571, 317–328. Springer, 2005.

[Dimopoulos and Torres, 1996] Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoret. Comput. Sci.*, 170(1-2):209–244, 1996.

[Downey and Fellows, 1999] R. G. Downey and M. R. Fellows. *Parameterized Complexity.* Springer, 1999.

[Dung, 1995] Ph. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.

[Dunne and Bench-Capon, 2001] P. E. Dunne and T. J. M. Bench-Capon. Complexity and combinatorial properties of argument systems. Technical report, University of Liverpool, 2001.

[Dunne and Bench-Capon, 2002] P. E. Dunne and T. J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141(1-2):187–203, 2002.

[Dunne, 2007] P. E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171(10-15):701–729, 2007.

[Dvořák and Woltran, 2010] W. Dvořák and S. Woltran. On the intertranslatability of argumentation semantics. In *Conference on Thirty Years of Nonmonotonic Reasoning (NonMon@30)*, Lexington, KY, USA, 2010.

[Dvořák *et al.*, 2010] W. Dvořák, S. Szeider, and S. Woltran. Reasoning in argumentation frameworks of bounded clique-width. In P. Baroni, F. Cerutti, M. Giacomin, and G. R. Simari, eds., *COMMA 2010*, volume 216 of *Frontiers in Artificial Intelligence and Applications*, pages 219–230. IOS, 2010.

[Fichte and Szeider, 2011] J. K. Fichte and S. Szeider. Backdoors to tractable answer-set programming. In *IJCAI 2011*, 2011.

[Flum and Grohe, 2006] J. Flum and M. Grohe. *Parameterized Complexity Theory*, volume XIV of *Texts in Theoretical Computer Science. An EATCS Series.* Springer Verlag, Berlin, 2006.

[Garey and Johnson, 1979] M. R. Garey and D. R. Johnson. *Computers and Intractability.* W. H. Freeman and Company, New York, San Francisco, 1979.

[Gottlob and Szeider, 2006] G. Gottlob and S. Szeider. Fixed-parameter algorithms for artificial intelligence, constraint satisfaction, and database problems. *The Computer Journal*, 51(3):303–325, 2006.

[Modgil and Caminada, 2009] S. Modgil and M. Caminada. Proof theories and algorithms for abstract argumentation frameworks. In I. Rahwan and G. Simari, eds., *Argumentation in Artificial Intelligence*, pages 105–132. Springer, 2009.

[Niedermeier, 2006] R. Niedermeier. *Invitation to Fixed-Parameter Algorithms.* Oxford Lecture Series in Mathematics and its Applications. Oxford University Press, Oxford, 2006.

[Parsons *et al.*, 2003] S. Parsons, M. Wooldridge, and L. Amgoud. Properties and complexity of some formal inter-agent dialogues. *J. Logic Comput.*, 13(3):347–376, 2003.

[Pollock, 1992] J. L. Pollock. How to reason defeasibly. *Artificial Intelligence*, 57(1):1–42, 1992.

[Rahwan and Simari, 2009] I. Rahwan and G. R. Simari, eds. *Argumentation in Artificial Intelligence.* Springer, 2009.

[Reed *et al.*, 2004] B. Reed, K. Smith, and A. Vetta. Finding odd cycle transversals. *Oper. Res. Lett.*, 32(4):299–301, 2004.

[Robertson *et al.*, 1999] N. Robertson, P. D. Seymour, and R. Thomas. Permanents, Pfaffian orientations, and even directed circuits. *Ann. of Math. (2)*, 150(3):929–975, 1999.

[Samer and Szeider, 2009] M. Samer and S. Szeider. Fixed-parameter tractability. In A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors, *Handbook of Satisfiability*, chapter 13, pages 425–454. IOS Press, 2009.

[Samer and Szeider, 2009a] M. Samer and S. Szeider. Backdoor sets of quantified Boolean formulas. *Journal of Automated Reasoning*, 42(1):77–97, 2009.

[Williams *et al.*, 2003] R. Williams, C. Gomes, and B. Selman. Backdoors to typical case complexity. In G. Gottlob and T. Walsh, eds, *IJCAI 2003*, pages 1173–1178, 2003.